PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant:	ASCHOFF et al.	Examiner:	TSAL, S. J.
Serial No.:	10/672,423	Group Art Unit:	2186
Filed:	September 26, 2003	Docket No.:	**SJO920030019US1** (IBMS.069-0510)
Title:	METHOD, APPARATUS AND PROGRAM STORAGE DEVICE FOR PROVIDING AUTOMATIC PERFORMANCE OPTIMIZATION OF VIRTUALIZED STORAGE ALLOCATION WITHIN A NETWORK OF STORAGE ELEMENTS		

SUPPLEMENTAL APPEAL BRIEF

(Correcting Listing of Claims in Non-Compliant

Appeal Brief of October 24, 2006)

Mail Stop Appeal
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This is an Appeal Brief submitted pursuant to 37 C.F.R. § 41.37 for the above-referenced patent application. Please charge Deposit Account No. 09-0466 (**SJO920030019US1**) in the amount of $500.00 for this brief in support of appeal as indicated in 37 C.F.R. § 41.20(b)(2).

I. Real Party In Interest

The real party in interest is International Business Machines Corporation, having a place of business at New Orchard Road, Armonk, New York 10504. This application is assigned to International Business Machines Corporation.

II. Related Appeals And Interferences

Appellant is unaware of any related appeals, interferences or judicial proceedings.

III. Status Of Claims

Claims 1-40 are rejected. Claims 41-43 were canceled in the prevsious response dated June 2, 2006. Claims 1-40 are presented for appeal and may be found in the attached Appendix of Appealed Claims in their present form.

IV. Status Of Amendments

No amendments to the claims were made subsequent to the final rejection of Appellants' application.

V. Summary Of Claimed Subject Matter

Appellant's invention is a method, apparatus and program storage device for providing automatic performance optimization of virtualized storage allocation within a network of storage elements.

Claim 1 is directed to an administration device (Fig. 2, 270, page 10, line 5) for providing automatic performance optimization of virtualized storage allocation within a network of storage elements. The administration device includes memory (Fig. 2, 292, page 18, line 10) for storing data thereon. The administration device also includes a processor (Fig. 2, 294, page 18, line 13) configured for receiving from a user a request for storage of data (Fig. 6, 610, page 17, lines 7-8) and for determining workload requirements of the user making the request (Fig. 6, 620, page 17, lines 7-8; Fig. 2, 272, page 13, lines 20-22). The processor is also configured for analyzing system parameters including performance characteristics of storage volumes within the network (Fig. 6, 630, page 17, lines 9-10) and for providing storage to meet the workload requirements of the user determined by the processor and to meet competing workload requirements based on the analysis of the system parameters (Fig. 6, 640, page 17, lines 10-12).

Claim 13 is directed to a network storage system (Fig. 2, 200, page 9, line 13). The network storage system includes a plurality of storage devices (Fig. 2, 230, page 9, line 14), a plurality of servers (Fig. 2, 202, page 9, line 14) coupled to the plurality of storage devices via network interconnect (Fig. 2, 204, page 9, lines 14-15) and an administration device (Fig. 2, 270, page 10, line 5), coupled to at least the plurality of storage devices, for providing automatic performance optimization of virtualized storage allocation within a network of storage elements. The administration device further includes memory (Fig. 2, 292, page 18,

line 10) for storing data thereon and a processor (Fig. 2, 294, page 18, line 13). The processor is configured for receiving from a user a request for storage of data (Fig. 6, 610, page 17, lines 7-8) and for determining workload requirements of the user making the request (Fig. 6, 620, page 17, lines 7-8; Fig. 2, 272, page 13, lines 20-22). The processor is also configured for analyzing system parameters including performance characteristics of storage volumes within the network (Fig. 6, 630, page 17, lines 9-10) and for providing storage to meet the workload requirements of the user determined by the processor and to meet competing workload requirements based on the analysis of the system parameters (Fig. 6, 640, page 17, lines 10-12).

Claim 25 is directed to a method (Fig. 6, 600, page 17, lines 5-7) for providing automatic performance optimization of virtualized storage allocation within a network of storage elements. The method includes receiving from a user a request for storage of data (Fig. 6, 610, page 17, lines 7-8), determining workload requirements of the user making the request (Fig. 6, 620, page 17, lines 7-8; Fig. 2, 272, page 13, lines 20-22), analyzing system parameters including performance characteristics of storage volumes within the network (Fig. 6, 630, page 17, lines 9-10) and providing storage to meet the determined workload requirements of the user and to meet competing workload requirements based on the analysis of the system parameters (Fig. 6, 640, page 17, lines 10-12).

Claim 37 is directed to a program storage device (Fig. 2, 288, page 18, lines 7-9) readable by a computer that tangibly embodies one or more programs of instructions (Fig. 2, 290, page 18, lines 12-15) executable by the computer to perform a method for providing automatic performance optimization of virtualized storage allocation within a network of storage elements. The method includes receiving from a user a request for storage of data

(Fig. 6, 610, page 17, lines 7-8), determining workload requirements of the user making the request (Fig. 6, 620, page 17, lines 7-8; Fig. 2, 272, page 13, lines 20-22), analyzing system parameters including performance characteristics of storage volumes within the network (Fig. 6, 630, page 17, lines 9-10) and providing storage to meet the determined workload requirements of the user and to meet competing workload requirements based on the analysis of the system parameters (Fig. 6, 640, page 17, lines 10-12).

Claim 39 is directed to an administration device (Fig. 2, 270, page 10, line 5) for providing automatic performance optimization of virtualized storage allocation within a network of storage elements. The administration device includes means for storing data thereon (Fig. 2, 292, page 18, line 10). The administration device also includes means (Fig. 2, 294, page 18, line 13) configured for receiving from a user a request for storage of data (Fig. 6, 610, page 17, lines 7-8), for determining workload requirements of the user making the request (Fig. 6, 620, page 17, lines 7-8; Fig. 2, 272, page 13, lines 20-22), for analyzing system parameters including performance characteristics of storage volumes within the network (Fig. 6, 630, page 17, lines 9-10) and for providing storage to meet the workload requirements of the user and to meet competing workload requirements based on the analysis of the system parameters (Fig. 6, 640, page 17, lines 10-12).

Claim 40 is directed to a network storage system (Fig. 2, 200, page 9, line 13). The network storage system includes first means for providing storage (Fig. 2, 230, page 9, line 14) and means for providing access to the means for providing storage (Fig. 2, 202, page 9, line 14). The network storage system also includes means (Fig. 2, 270, page 10, line 5), coupled to at least the plurality of storage devices, for providing automatic performance optimization of virtualized storage allocation within a network of storage elements. The

administration device further includes second means (Fig. 2, 292, page 18, line 10) for storing data thereon and means (Fig. 2, 294, page 18, line 13) configured for receiving from a user a request for storage of data (Fig. 6, 610, page 17, lines 7-8), for determining workload requirements of the user making the request (Fig. 6, 620, page 17, lines 7-8; Fig. 2, 272, page 13, lines 20-22), for analyzing system parameters including performance characteristics of storage volumes within the network (Fig. 6, 630, page 17, lines 9-10) and for providing storage to meet the workload requirements of the user and to meet competing workload requirements based on the analysis of the system parameters (Fig. 6, 640, page 17, lines 10-12).

VI. Grounds Of Rejections To Be Reviewed On Appeal

Appellant has attempted to comply with new rule 37 C.F.R. § 41.37(c) by providing the Office Action's grounds of rejection verbatim, followed by an argument section corresponding thereto.

A. Claims 1-9, 12-21, 24-33 and 36-40 were rejected under 35 U.S.C. § 102(e) as being anticipated by Zahavi.

B. Claims 10-11, 22-23, 34-35 and 41-42 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Zahavi in view of Lee et al.

VII. Argument

A. REJECTION OF CLAIMS 1-9, 12-21, 24-33 AND 36-40 UNDER 35 U.S.C. 102(e) AS BEING ANTICIPATED BY ZAHAVI

1. INDEPENDENT CLAIMS 1, 13, 39 AND 40 ARE PATENTABLE OVER ZAHAVI

i. ZAHAVI FAILS TO TEACH, DISCLOSE OR SUGGEST A PROCESSOR CONFIGURED FOR DETERMINING WORKLOAD REQUIREMENTS OF A USER MAKING A REQUEST FOR STORAGE OF DATA

Regarding independent claims 1, 13, 25, 37, 39 and 40, a processor is configured for determining workload requirements of a user. A user issues a request for storage of data and the processor determines the workload requirements of the user based upon the request.

In contrast, Zahavi explicitly requires the user to identify workload requirements, e.g., what percent of the work is Random-Read Hit, Random-Read Miss, Sequential Read and Write (see column 8, lines 1-5). The user is also asked to provide other information, e.g., type and size of traffic (column 7, lines 13-15), adjust space requirements (column 8, lines 64-65), selection of the best choice for data storage system (column 10, lines 8-9), etc.

Thus, according to Zahavi, the user determines the workload requirements rather than providing a processor configured for determining workload requirements of a user.

Therefore, Appellants respectfully submit that independent claims 1, 13, 25, 37, 39 and 40 are patentable over Zahavi.

ii. ZAHAVI FAILS TO TEACH, DISCLOSE OR SUGGEST A PROCESSOR CONFIGURED FOR PROVIDING STORAGE TO MEET THE WORKLOAD REQUIREMENTS OF THE USER DETERMINED BY THE PROCESSOR AND TO MEET COMPETING WORKLOAD REQUIREMENTS BASED ON THE ANALYSIS OF THE SYSTEM PARAMETERS

According to independent claims 1, 13, 25, 37, 39 and 40, the processor is also configured for providing storage to meet the workload requirements of the user determined by the processor.

As indicated above, Zahavi requires the user to identify workload requirements. The user is also asked to provide other information. Accordingly, Zahavi does not even suggest a processor that is also configured for providing storage to meet the workload requirements of the user determined by the processor.

Therefore, Appellants respectfully submit that independent claims 1, 13, 25, 37, 39 and 40 are patentable over Zahavi.

2. INDEPENDENT CLAIMS 25 AND 37 ARE PATENTABLE OVER ZAHAVI

i. ZAHAVI FAILS TO TEACH, DISCLOSE OR SUGGEST A METHOD FOR PROVIDING AUTOMATIC PERFORMANCE OPTIMIZATION OF VIRTUALIZED STORAGE ALLOCATION WITHIN A NETWORK OF STORAGE ELEMENTS BY DETERMINING WORKLOAD REQUIREMENTS OF A USER MAKING A REQUEST FOR STORAGE OF DATA

As required by independent claims 25 and37, the method requires providing automatic performance optimization of virtualized storage allocation within a network of storage elements by determining workload requirements of a user making a request for storage of data.

In contrast, Zahavi explicitly requires the user to identify workload requirements, e.g., what percent of the work is Random-Read Hit, Random-Read Miss, Sequential Read and Write (see column 8, lines 1-5). The user is also asked to provide other information, e.g., type and size of traffic (column 7, lines 13-15), adjust space requirements (column 8, lines 64-65), selection of the best choice for data storage system (column 10, lines 8-9), etc.

Thus, Zahavi discloses that the user determines the workload requirements rather than disclosing a method for providing automatic performance optimization of virtualized storage allocation within a network of storage elements by determining workload requirements of a user making a request for storage of data. The system of Zahavi is not automatic, but rather a manual system, if the user must provide input to identify the workload requirements.

Therefore, Appellants respectfully submit that independent claims 25 and 37 are patentable over Zahavi.

2. DEPENDENT CLAIMS 6 AND 18 ARE PATENTABLE OVER ZAHAVI

Dependent claims 6 and 18 recite a processor that determines workload attributes of the user and desired levels of performance.

As described above with reference to independent claims 1, 13, 39 and 40, Zahavi relies upon user input for the workload attributes. However, Zahavi also relies upon the user to supply an indication of a desired level of performance. For example, Zahavi discloses that screen area 514 "includes an area to enter a Performance Zone Value" (column 6, lines 47-48). Accordingly, Zahavi fails to suggest that a processor determines workload attributes of the user and desired levels of performance.

Therefore, Appellants respectfully submit that dependent claims 6 and 18 are patentable over Zahavi.

3. DEPENDENT CLAIMS 7 AND 19 ARE PATENTABLE OVER ZAHAVI

Dependent claims 7 and 19 recite a processor that is configured for determining workload requirements of the user by using canned workload descriptions that are based on characterizations of user environments across various industries and applications.

In contrast, Zahavi provides a Workload Library. However, according to Zahavi the user determines workload requirements using knowledge of the application transactions or a Workload Library.

Accordingly, Zahavi fails to disclose, teach or suggest a processor that is configured for determining workload requirements of the user by using canned workload descriptions that are based on characterizations of user environments across various industries and applications.

Therefore, Appellants respectfully submit that dependent claims 7 and 19 are patentable over Zahavi.

4. DEPENDENT CLAIMS 8 AND 20 ARE PATENTABLE OVER ZAHAVI

Dependent claims 8 and 20 recite a processor that is configured for determining workload requirements of the user by automatically creating workload requirements based on observations of storage access patterns of a user.

Zahavi fails to even mention the use of observations of storage access patterns of a user. Zahavi does disclose providing a display of the activity patterns of each affinity group so that the validity of the selection can be visualized. However, this involves the mere presentation of a visualization of affinity group activity patterns and does not equate to providing a processor that is configured for determining workload requirements of the user by automatically creating workload requirements based on observations of storage access patterns of a user.

First, the display relates to activity patterns associated with correlated pairs of storage devices and not the storage access patterns of a user.

Further, the displaying of activity patterns associated with correlated pairs of storage devices does not relate to determining workload requirements of the user by automatically creating workload requirements based on observations of storage access patterns of a user.

Accordingly, Zahavi fails to disclose, teach or suggest a processor that is configured for determining workload requirements of the user by automatically creating workload requirements based on observations of storage access patterns of a user.

Therefore, Appellants respectfully submit that dependent claims 8 and 20 are patentable over Zahavi.

5. DEPENDENT CLAIMS 9 AND 21 ARE PATENTABLE OVER ZAHAVI

Dependent claims 9 and 21 recite a processor that is configured for determining workload requirements of the user by using intelligent software components that analyze workload descriptions for an application of the user.

In contrast, Zahavi merely mentions analyzing I/O activity. Accordingly, Zahavi fails to disclose, teach or suggest analyzing workload descriptions for an application of the user.

Thus, Zahavi fails to disclose, teach or suggest a processor that is configured for determining workload requirements of the user by using intelligent software components that analyze workload descriptions for an application of the user.

Therefore, Appellants respectfully submit that dependent claims 9 and 21 are patentable over Zahavi.

B. REJECTION OF CLAIMS 10-11, 22-23, 34-35 AND 41-42 UNDER 35 U.S.C. § 103(a) AS BEING UNPATENTABLE OVER ZAHAVI IN VIEW OF LEE ET AL.

Dependent claims 10-11, 22-23 and 34-35 are also patentable over the cited reference, because they incorporate all of the limitations of the corresponding independent claim 1, 13 and 25, respectively. Claims 41-42 have been canceled. Therefore, Applicants respectfully submit that dependent claims 10-11, 22-23, 34-35 and 41-42 are patentable over Zahavi and Lee et al., and request that the rejections to the dependent claims be withdrawn.

C. Conclusion

In view of the above, Appellant submits that the rejections are improper, the claimed invention is patentable, and that the rejections of claims 1-40 should be reversed. Appellant respectfully requests reversal of the rejections as applied to the appealed claims and allowance of the entire application.

Respectfully submitted,

Chambliss, Bahner and Stophel
1000 Tallan Building
Two Union Square
Chattanooga, TN 37402
423-757-0264

By: 

Name: David W. Lynch
Reg. No.: 36,204

VIII. Claims Appendix

1. (Previously Presented) An administration device for providing automatic performance optimization of virtualized storage allocation within a network of storage elements, comprising:

memory for storing data thereon; and

a processor configured for receiving from a user a request for storage of data, for determining workload requirements of the user making the request, for analyzing system parameters including performance characteristics of storage volumes within the network and for providing storage to meet the workload requirements of the user determined by the processor and to meet competing workload requirements based on the analysis of the system parameters.

2. (Original) The administration device of claim 1, wherein the processor provides storage to meet the workload requirements based on the workload requirements of the user and storage requirements for the data.

3. (Original) The administration device of claim 1, wherein the processor provides storage to meet the workload requirements by selecting storage locations that meet performance and space requirements of the request.

4. (Original) The administration device of claim 3, wherein the processor selects storage locations that meet the performance and space requirements through analysis of the request for storage.

5. (Original) The administration device of claim 3, wherein the processor selects storage locations that meet the performance and space requirements through a storage policy mechanism.

6. (Original) The administration device of claim 1, wherein the processor determines workload attributes of the user and desired levels of performance, retains the latest information about the available capacity within the network of storage elements, determines performance characteristics of individual storage devices at different locations within the network as a function of the workload requirements of the user, and determines a presence and attributes of competing workloads sharing the storage devices over extended periods of time.

7. (Previously Presented) The administration device of claim 1, wherein the processor is configured for determining workload requirements of the user by using canned workload descriptions that are based on characterizations of user environments across various industries and applications.

8. (Previously Presented) The administration device of claim 1, wherein the processor is configured for determining workload requirements of the user by automatically creating workload requirements based on observations of storage access patterns of a user.

9. (Previously Presented) The administration device of claim 1, wherein the processor is configured for determining workload requirements of the user by using intelligent software components that analyze workload descriptions for an application of the user.

10. (Original) The administration device of claim 1, wherein the processor is configured for accessing a virtualization engine and volume managers to stripe data within a virtual disk across managed storage devices.

11. (Original) The administration device of claim 1, wherein the processor is configured for determining how to relocate virtual disks to meet a desired level of performance.

12. (Original) The administration device of claim 1, wherein the processor is configured for performing a calibration process to discover the performance capabilities of the underlying storage devices.

13. (Previously Presented) A network storage system, comprising:

a plurality of storage devices;

a plurality of servers coupled to the plurality of storage devices via network interconnect; and

an administration device, coupled to at least the plurality of storage devices, for providing automatic performance optimization of virtualized storage allocation within a network of storage elements, wherein the administration device further comprises:

memory for storing data thereon; and

a processor configured for receiving from a user a request for storage of data, for determining workload requirements of the user making the request, for analyzing system parameters including performance characteristics of storage volumes within the network and for providing storage to meet the workload requirements of the user determined by the processor and to meet competing workload requirements based on the analysis of the system parameters.

14. (Original) The network storage system of claim 13, wherein the processor provides storage to meet the workload requirements based on the workload requirements of the user and storage requirements for the data.

15. (Original) The network storage system of claim 13, wherein the processor provides storage to meet the workload requirements by selecting storage locations that meet performance and space requirements of the request.

16. (Original) The network storage system of claim 15, wherein the processor selects storage locations that meet the performance and space requirements through analysis of the request for storage.

17. (Original) The network storage system of claim 15, wherein the processor selects storage locations that meet the performance and space requirements through a storage policy mechanism.

18. (Original) The network storage system of claim 13, wherein the processor determines workload attributes of the user and desired levels of performance, retains the latest information about the available capacity within the network of storage elements, determines performance characteristics of individual storage devices at different locations within the network as a function of the workload requirements of the user, and determines a presence and attributes of competing workloads sharing the storage devices over extended periods of time.

19. (Previously Presented) The network storage system of claim 13, wherein the processor is configured for determining workload requirements of the user by using canned workload descriptions that are based on characterizations of user environments across various industries and applications.

20. (Previously Presented) The network storage system of claim 13, wherein the processor is configured for determining workload requirements of the user by automatically creating workload requirements based on observations of storage access patterns of a user.

21. (Previously Presented) The network storage system of claim 13, wherein the processor is configured for determining workload requirements of the user by using intelligent software components that analyze workload descriptions for an application of the user.

22. (Original) The network storage system of claim 13, wherein the processor is configured for accessing a virtualization engine and volume managers to stripe data within a virtual disk across managed storage devices.

23. (Original) The network storage system of claim 13, wherein the processor is configured for determining how to relocate virtual disks to meet a desired level of performance.

24. (Original) The network storage system of claim 13, wherein the processor is configured for performing a calibration process to discover the performance capabilities of the underlying storage devices.

25. (Previously Presented) A method for providing automatic performance optimization of virtualized storage allocation within a network of storage elements, comprising:

receiving from a user a request for storage of data;

determining workload requirements of the user making the request;

analyzing system parameters including performance characteristics of storage volumes within the network; and

providing storage to meet the determined workload requirements of the user and to meet competing workload requirements based on the analysis of the system parameters.

26. (Original) The method of claim 25, wherein the providing storage to meet the workload requirements of the user is further based on the workload requirements of the user and storage requirements for the data.

27. (Original) The method of claim 25, wherein the providing storage to meet the workload requirements of the user further comprises selecting storage locations that meet performance and space requirements of the request.

28. (Original) The method of claim 27, wherein the selecting storage locations that meet the performance and space requirements are provided with the request for storage.

29. (Original) The method of claim 27, wherein the selecting storage locations that meet the performance and space requirements are provided through a storage policy mechanism.

30. (Original) The method of claim 25, wherein the analyzing system parameters further comprises determining the workload attributes of the user and desired levels of performance, retaining the latest information about the available capacity within the network of storage elements, determining performance characteristics of the storage devices at different locations within the network as a function of the workload requirements of the user, and determining a presence and attributes of competing workloads sharing the storage devices over extended periods of time.

31. (Previously Presented) The method of claim 25, wherein the determining workload requirements of the user making the request further comprises using canned workload descriptions that are based on characterizations of user environments across various industries and applications.

32. (Previously Presented) The method of claim 25, wherein the determining workload requirements of the user making the request further comprises automatically creating workload requirements based on observations of storage access patterns of a user.

33. (Previously Presented) The method of claim 25, wherein the determining workload requirements of the user making the request further comprises using intelligent software components that analyze workload descriptions for an application of the user.

34. (Original) The method of claim 25 further comprising accessing a virtualization engine and volume managers to stripe data within a virtual disk across managed storage devices.

35. (Original) The method of claim 34, wherein the striping data further comprises determining how to relocate virtual disks to meet a desired level of performance.

36. (Original) The method of claim 25, further comprising performing a calibration process to discover the performance capabilities of the underlying storage devices.

37. (Previously Presented) A program storage device readable by a computer, the program storage device tangibly embodying one or more programs of instructions executable by the computer to perform a method for providing automatic performance optimization of virtualized storage allocation within a network of storage elements, the method comprising:

receiving from a user a request for storage of data;

determining workload requirements of the user making the request;

analyzing system parameters including performance characteristics of storage volumes within the network; and

providing storage to meet the determined workload requirements of the user and to meet competing workload requirements based on the analysis of the system parameters.

38. (Original) The program storage device of claim 19, wherein the analyzing system parameters further comprises determining workload attributes of the user and desired levels of performance, retaining the latest information about the available capacity within the network of storage elements, determining performance characteristics of storage devices at different locations within the network as a function of the workload requirements of the user, and determining a presence and attributes of competing workloads sharing the storage devices over extended periods of time.

39. (Previously Presented) An administration device for providing automatic performance optimization of virtualized storage allocation within a network of storage elements, comprising:

means for storing data thereon; and

means configured for receiving from a user a request for storage of data, for determining workload requirements of the user making the request, for analyzing system parameters including performance characteristics of storage volumes within the network and for providing storage to meet the workload requirements of the user and to meet competing workload requirements based on the analysis of the system parameters.

40. (Previously Presented) A network storage system, comprising:

first means for providing storage;

means for providing access to the means for providing storage; and

means, coupled to at least the plurality of storage devices, for providing automatic performance optimization of virtualized storage allocation within a network of storage elements, wherein the administration device further comprises:

second means for storing data thereon; and

means configured for receiving from a user a request for storage of data, for determining workload requirements of the user making the request, for analyzing system parameters including performance characteristics of storage volumes within the network and for providing storage to meet the workload requirements of the user and to meet competing workload requirements based on the analysis of the system parameters.

41-43. (Canceled)

IX. Evidence Appendix

Appellant is unaware of any evidence submitted in this application pursuant to 37 C.F.R. §§ 1.130, 1.131, and 1.132.

X. Related Proceedings Appendix

As stated in Section II above, Appellant is unaware of any related appeals, interferences or judicial proceedings.